

	Type	Hits	Search Text	DBs
1	BRS	160	"data cube" and dimension	USPAT; EPO; JPO; IBM_TDB
2	BRS	30	("data cube" and dimension) and "data mining"	USPAT; EPO; JPO; IBM_TDB
3	BRS	6	((("data cube" and dimension) and "data mining")) and star	USPAT; EPO; JPO; IBM_TDB
4	BRS	0	(edit near2 "data cube") and (aggregate or aggregation) and star	USPAT; EPO; JPO; IBM_TDB
5	BRS	0	edit near2 "data cube"	USPAT; EPO; JPO; IBM_TDB
6	BRS	1	(edit or modify or change) near2 "data cube"	USPAT; EPO; JPO; IBM_TDB
7	BRS	24	((("data cube" and dimension) and "data mining")) and (aggregate or aggregation)	USPAT; EPO; JPO; IBM_TDB
8	BRS	6	((("data cube" and dimension) and "data mining")) and (aggregate or aggregation)) and star	USPAT; EPO; JPO; IBM_TDB
9	BRS	10	((("data cube" and dimension) and "data mining")) and (aggregate or aggregation)) and schema\$	USPAT; EPO; JPO; IBM_TDB
10	BRS	0	((("data cube" and dimension) and "data mining")) and (aggregate or aggregation)) and anonym\$	USPAT; EPO; JPO; IBM_TDB
11	BRS	0	((("data cube" and dimension) and "data mining")) and anonym\$	USPAT; EPO; JPO; IBM_TDB
12	BRS	2	("data cube" and dimension) and anonym\$	USPAT; EPO; JPO; IBM_TDB
13	BRS	2	("data cube" and dimension) and anonymous	USPAT; EPO; JPO; IBM_TDB
14	BRS	0	("data cube" and dimension) and anonymity	USPAT; EPO; JPO; IBM_TDB
15	BRS	5	((edit or modify or change) near2 dimension) and "data cube"	USPAT; EPO; JPO; IBM_TDB

	Type	Hits	Search Text	DBs
16	BRS	44	"data cube".ab,ti,clm.	USPAT; EPO; JPO; IBM_TDB
17	BRS	12	"data cube".ab,ti,clm. and (query near4 ("data cube" or dimension))	USPAT; EPO; JPO; IBM_TDB
18	BRS	2	("data cube".ab,ti,clm. and (query near4 ("data cube" or dimension))) and star	USPAT; EPO; JPO; IBM_TDB
19	BRS	0	6633882.pn. and star	USPAT; EPO; JPO; IBM_TDB
20	BRS	0	6633882.pn. and schema	USPAT; EPO; JPO; IBM_TDB
21	BRS	1	6721749.pn. and star	USPAT; EPO; JPO; IBM_TDB
22	BRS	0	6633882.pn. and schemas	USPAT; EPO; JPO; IBM_TDB
23	BRS	1	6633882.pn. and schem\$	USPAT; EPO; JPO; IBM_TDB



US Patent &amp; Trademark Office

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide

"data cube" and star and sql

SEARCH


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)
Terms used data cube and star and sql

Found 588 of 138,663

Sort results by

relevance

[Save results to a Binder](#)[Try an Advanced Search](#)[Try this search in The ACM Guide](#)

Display results

expanded form

[Search Tips](#)☐ Open results in a new window

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale ☐ ☐ ☐ ☐ ☐**1 [Query processing: Implementing operations to navigate semantic star schemas](#)**

Alberto Abelló, José Samos, Fèlix Saltor

November 2003 **Proceedings of the 6th ACM international workshop on Data warehousing and OLAP**Full text available: [pdf\(193.82 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

In the last years, lots of work have been devoted to multidimensional modeling, star shape schemas and OLAP operations. However, "drill-across" has not captured as much attention as other operations. This operation allows to change the subject of analysis keeping the same analysis space we were using to analyze another subject. It is assumed that this can be done if both subjects share exactly the same analysis dimensions. In this paper, besides the implementation of an algebraic set of operatio ...

**Keywords:** OLAP operations, SQL, drill-across, semantic relationships, star schema**2 [Knowledge discovery in data warehouses](#)**

Themistoklis Palpanas

September 2000 **ACM SIGMOD Record**, Volume 29 Issue 3Full text available: [pdf\(240.77 KB\)](#) Additional Information: [full citation](#), [abstract](#), [index terms](#)

As the size of data warehouses increase to several hundreds of gigabytes or terabytes, the need for methods and tools that will automate the process of knowledge extraction, or guide the user to subsets of the dataset that are of particular interest, is becoming prominent. In this survey paper we explore the problem of identifying and extracting interesting knowledge from large collections of data residing in data warehouses, by using data mining techniques. Such techniques have the ability to i ...

**3 [Poster papers - short papers: A visual interface technique for exploring OLAP data with coordinated dimension hierarchies](#)**

Mark Sifer

November 2003 **Proceedings of the twelfth international conference on Information and knowledge management**Full text available: [pdf\(272.82 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Multi-dimensional data occurs in many domains while a wide variety of text based and visual interfaces for querying such data exists. But many of these interfaces are not applicable to OLAP, as they do not support use of dimension hierarchies for selection and aggregation. We

introduce an interface technique which supports visual querying of OLAP data, that has been implemented in the SGViewer tool. It is based on a data graph rather than a data cube representation of the data. Our interface pre ...

**Keywords:** OLAP, data exploration, hierarchies, interface

#### 4 A survey of logical models for OLAP databases

Panos Vassiliadis, Timos Sellis

December 1999 **ACM SIGMOD Record**, Volume 28 Issue 4

Full text available:  [pdf\(604.36 KB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

In this paper, we present different proposals for multidimensional data cubes, which are the basic logical model for OLAP applications. We have grouped the work in the field in two categories: commercial tools (presented along with terminology and standards) and academic efforts. We further divide the academic efforts in two subcategories: the relational model extensions and the cube-oriented approaches. Finally, we attempt a comparative analysis of the various efforts.

#### 5 CubiST: a new algorithm for improving the performance of ad-hoc OLAP queries

Lixin Fu, Joachim Hammer

November 2000 **Proceedings of the 3rd ACM international workshop on Data warehousing and OLAP**


Full text available:  [pdf\(296.08 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

**Keywords:** OLAP, data cube, data warehouse, index structure, query optimization, query processing

#### 6 On relationships offering new drill-across possibilities

Alberto Abelló, José Samos, Fèlix Saltor

November 2002 **Proceedings of the 5th ACM international workshop on Data Warehousing and OLAP**

Full text available:  [pdf\(236.84 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

OLAP tools divide concepts based on whether they are used as analysis dimensions, or are the fact subject of analysis, which gives rise to star shape schemas. Operations are always provided to navigate inside such star schemas. However, the navigation among different stars is usually overlooked. This paper studies different kinds of Object-Oriented conceptual relationships (part of UML standard) between stars (namely *Derivation*, *Generalization*, *Association*, and *Flow*) ...

**Keywords:** UML, drill-across, multidimensional design, semantics

#### 7 Maintenance of data cubes and summary tables in a warehouse

Inderpal Singh Mumick, Dallan Quass, Barinderpal Singh Mumick

June 1997 **ACM SIGMOD Record**, **Proceedings of the 1997 ACM SIGMOD international conference on Management of data**, Volume 26 Issue 2

Full text available:  [pdf\(1.58 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Data warehouses contain large amounts of information, often collected from a variety of independent sources. Decision-support functions in a warehouse, such as on-line analytical processing (OLAP), involve hundreds of complex aggregate queries over large volumes of

data. It is not feasible to compute these queries by scanning the data sets each time. Warehouse applications therefore build a large number of summary tables, or materialized aggregate views, to ...

#### 8 Requirement-based data cube schema design

David W. Cheung, Bo Zhou, Ben Kao, Hongjun Lu, Tak Wah Lam, Hing Fung Ting

November 1999 **Proceedings of the eighth international conference on Information and knowledge management**

Full text available:  [pdf\(1.02 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

On-line analytical processing (OLAP) requires efficient processing of complex decision support queries over very large databases. It is well accepted that pre-computed data cubes can help reduce the response time of such queries dramatically. A very important design issue of an efficient OLAP system is therefore the choice of the right data cubes to materialize. We call this problem the data cube schema design problem. In this paper we show that the problem of finding an op ...

**Keywords:** DSS, OLAP, data cube schema design, data cubes

#### 9 Optimizing multiple dimensional queries simultaneously in multidimensional databases

Weifa Liang, Maria E. Orlowska, Jeffrey X. Yu

February 2000 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 8 Issue 3-4

Full text available:  [pdf\(269.57 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

Some significant progress related to multidimensional data analysis has been achieved in the past few years, including the design of fast algorithms for computing data cubes, selecting some precomputed group-bys to materialize, and designing efficient storage structures for multidimensional data. However, little work has been carried out on multidimensional query optimization issues. Particularly the response time (or evaluation cost) for answering several related dimensional queries simultaneous ...

**Keywords:** Data warehousing, MDDBs, Multiple dimensional query optimization, OLAP, Query modeling

#### 10 A powerful and SQL-compatible data model and query language for OLAP

Dennis Pedersen, Karsten Riis, Torben Bach Pedersen

January 2002 **Australian Computer Science Communications , Proceedings of the thirteenth Australasian conference on Database technologies - Volume 5**, Volume 24 Issue 2

Full text available:  [pdf\(1.12 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In this paper we present the SQLM OLAP data model, formal algebra, and query language that, unlike current OLAP data models and languages, are both *powerful*, meaning that they support irregular dimension hierarchies, automatic aggregation of data, and correct aggregation of data, and *SQL-compatible*, allowing seamless integration with relational technology. We also consider the requirements to the data model posed by integration of OLAP data with external XML data. ...

**Keywords:** OLAP, data integration, data models, multidimensional databases, query languages

### 11 Heuristic optimization of OLAP queries in multidimensionally hierarchically clustered databases

Dimitri Theodoratos, Aris Tsois

November 2001 **Proceedings of the 4th ACM international workshop on Data warehousing and OLAP**

Full text available:  [pdf\(1.44 MB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

On-line analytical processing (OLAP) is a technology that encompasses applications requiring a multidimensional and hierarchical view of data. OLAP applications often require fast response time to complex grouping/aggregation queries on enormous quantities of data. Commercial relational database management systems use mainly multiple one-dimensional indexes to process OLAP queries that restrict multiple dimensions. However, in many cases, multidimensional access methods outperform one-dimensiona ...

### 12 An overview of data warehousing and OLAP technology

Surajit Chaudhuri, Umeshwar Dayal

March 1997 **ACM SIGMOD Record**, Volume 26 Issue 1

Full text available:  [pdf\(101.60 KB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

Data warehousing and on-line analytical processing (OLAP) are essential elements of decision support, which has increasingly become a focus of the database industry. Many commercial products and services are now available, and all of the principal database management system vendors now have offerings in these areas. Decision support places some rather different requirements on database technology compared to traditional on-line transaction processing applications. This paper provides an overview ...

### 13 Simultaneous optimization and evaluation of multiple dimensional queries

Yihong Zhao, Prasad M. Deshpande, Jeffrey F. Naughton, Amit Shukla

June 1998 **ACM SIGMOD Record , Proceedings of the 1998 ACM SIGMOD international conference on Management of data**, Volume 27 Issue 2


Full text available:  [pdf\(1.52 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Database researchers have made significant progress on several research issues related to multidimensional data analysis, including the development of fast cubing algorithms, efficient schemes for creating and maintaining precomputed group-bys, and the design of efficient storage structures for multidimensional data. However, to date there has been little or no work on multidimensional query optimization. Recently, Microsoft has proposed "OLE DB for OLAP" as a standard multidime ...

### 14 OLAP and statistical databases: similarities and differences

Arie Shoshani

May 1997 **Proceedings of the sixteenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems**

Full text available:  [pdf\(1.66 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

### 15 A case for dynamic view management

Yannis Kotidis, Nick Roussopoulos

December 2001 **ACM Transactions on Database Systems (TODS)**, Volume 26 Issue 4

Full text available:  [pdf\(892.57 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Materialized aggregate views represent a set of redundant entities in a data warehouse that are frequently used to accelerate On-Line Analytical Processing (OLAP). Due to the complex structure of the data warehouse and the different profiles of the users who submit queries,


there is need for tools that will automate and ease the view selection and management processes. In this article we present DynaMat, a system that manages dynamic collections of materialized aggregate views in a data warehouse ...

**Keywords:** Data cube, OLAP, data warehousing, materialized views

**16 DynaMat: a dynamic view management system for data warehouses**

Yannis Kotidis, Nick Roussopoulos

June 1999 **ACM SIGMOD Record , Proceedings of the 1999 ACM SIGMOD international conference on Management of data**, Volume 28 Issue 2

Full text available:  pdf(1.44 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Pre-computation and materialization of views with aggregate functions is a common technique in Data Warehouses. Due to the complex structure of the warehouse and the different profiles of the users who submit queries, there is need for tools that will automate the selection and management of the materialized data. In this paper we present DynaMat, a system that dynamically materializes information at multiple levels of granularity in order to match the demand (workload) but also takes into ...

**17 An alternative storage organization for ROLAP aggregate views based on cubetrees**

Yannis Kotidis, Nick Roussopoulos

June 1998 **ACM SIGMOD Record , Proceedings of the 1998 ACM SIGMOD international conference on Management of data**, Volume 27 Issue 2

Full text available:  pdf(1.19 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The Relational On-Line Analytical Processing (ROLAP) is emerging as the dominant approach in data warehousing with decision support applications. In order to enhance query performance, the ROLAP approach relies on selecting and materializing in summary tables appropriate subsets of aggregate views which are then engaged in speeding up OLAP queries. However, a straight forward relational storage implementation of materialized ROLAP views is immensely wasteful on storage and incredibly inadequate ...

**18 Query processing: Spatial hierarchy and OLAP-favored search in spatial data warehouse**

Fangyan Rao, Long Zhang, Xiu Lan Yu, Ying Li, Ying Chen

November 2003 **Proceedings of the 6th ACM international workshop on Data warehousing and OLAP**

Full text available:  pdf(197.48 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Data warehouse and Online Analytical Processing(OLAP) play a key role in business intelligent systems. With the increasing amount of spatial data stored in business database, how to utilize these spatial information to get insight into business data from the geo-spatial point of view is becoming an important issue of data warehouse and OLAP. However, traditional data warehouse and OLAP tools can not fully exploit spatial data in coordinates because multi-dimensional spatial data does not have im ...

**Keywords:** spatial OLAP, spatial data warehouse, spatial hierarchy

**19 Improved query performance with variant indexes**

Patrick O'Neil, Dallan Quass

June 1997 **ACM SIGMOD Record , Proceedings of the 1997 ACM SIGMOD international conference on Management of data**, Volume 26 Issue 2

Additional Information:

Full text available:  [pdf\(1.54 MB\)](#)

[full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The read-mostly environment of data warehousing makes it possible to use more complex indexes to speed up queries than in situations where concurrent updates are present. The current paper presents a short review of current indexing technology, including row-set representation by Bitmaps, and then introduces two approaches we call Bit-Sliced indexing and Projection indexing. A Projection index materializes all values of a column in RID order, and a Bit-Sliced index essentially takes an orth ...

## 20 Dealing with slow-evolving fact: a case study on inventory data warehousing

Chung-Min Chen, Munir Cochinwala, Elsa Yueh

November 1999 **Proceedings of the 2nd ACM international workshop on Data warehousing and OLAP**

Full text available:  [pdf\(941.25 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Data Warehousing for INventory management (DWIN) is a production project at Telcordia aimed at providing telecommunications service providers with decision support functions for inventory control and monitoring. In this paper, we report some interesting issues related to the design of the data warehouse. Specifically, we will discuss the issues of slow-evolving fact, transaction-oriented fact table, and large dimensions. We also propose the concept of virtual data cubes and ...

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.  
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)





US Patent &amp; Trademark Office

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide

edit data cube, dimension, star

SEARCH


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)
Terms used edit data cube dimension star

Found 44,188 of 138,663

Sort results by

relevance

Display results

expanded form

Save results to a Binder

Search Tips

☐ Open results in a new window
Try an [Advanced Search](#)Try this search in [The ACM Guide](#)

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale ☐ ☐ ☐ ☐ ☐

### 1 [Optimizing multiple dimensional queries simultaneously in multidimensional databases](#)

Weifa Liang, Maria E. Orlowska, Jeffrey X. Yu

February 2000 **The VLDB Journal — The International Journal on Very Large Data****Bases**, Volume 8 Issue 3-4Full text available: pdf(269.57 KB) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

Some significant progress related to multidimensional data analysis has been achieved in the past few years, including the design of fast algorithms for computing datacubes, selecting some precomputed group-bys to materialize, and designing efficient storage structures for multidimensional data. However, little work has been carried out on multidimensional query optimization issues. Particularly the response time (or evaluation cost) for answering several related dimensional queries simultaneous ...

**Keywords:** Data warehousing, MDDBs, Multiple dimensional query optimization, OLAP, Query modeling

### 2 [A method for developing dimensional data marts](#)

Tim Chenoweth, David Schuff, Robert St. Louis

December 2003 **Communications of the ACM**, Volume 46 Issue 12
 Full text available: pdf(115.16 KB) html(26.71 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Decision-oriented dimensional data marts are fundamentally different than transaction-oriented relational databases. A distinctive methodology and a different set of tools are required for their effective development.

### 3 [Deriving initial data warehouse structures from the conceptual data models of the underlying operational information systems](#)

Michael Boehnlein, Achim Ulbrich-vom Ende

November 1999 **Proceedings of the 2nd ACM international workshop on Data warehousing and OLAP**
 Full text available: pdf(1.40 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In recent years the construction of large scale data schemes for operational systems has been the major problem of conceptual data modeling for business needs. Multidimensional data structures used for decision support applications in data warehouses have rather

different requirements to data modeling techniques. In case of operational systems the data models are created from application specific requirements. The data models in data warehouses base on the analytical requirements of the use ...

**Keywords:** conceptual data model, data warehouse, decision support system, entity relationship model (ERM), snowflake scheme, star scheme, structured entity relationship model (SERM)

#### 4 What have we learnt from using real parallel machines to solve real problems?

G. C. Fox

January 1989 **Proceedings of the third conference on Hypercube concurrent computers and applications - Volume 2**

Full text available:  [pdf\(4.08 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We briefly review some key scientific and parallel processing issues in a selection of some 84 existing applications of parallel machines. We include the MIMD hypercube transputer array, BBN Butterfly, and the SIMD ICL DAP, Goodyear MPP and Connection Machine from Thinking Machines. We use a space-time analogy to classify problems and show how a division into synchronous, loosely synchronous and asynchronous problems is helpful. This classifies problems into those suitable for SIMD or MIMD ...

#### 5 Dynamic maintenance of multidimensional range data partitioning for parallel data processing

Junping Sun, William I. Grosky

November 1998 **Proceedings of the 1st ACM international workshop on Data warehousing and OLAP**

Full text available:  [pdf\(1.09 MB\)](#)

Additional Information: [full citation](#), [references](#), [index terms](#)

#### 6 Bottom-up computation of sparse and Iceberg CUBE

Kevin Beyer, Raghu Ramakrishnan

June 1999 **ACM SIGMOD Record , Proceedings of the 1999 ACM SIGMOD international conference on Management of data**, Volume 28 Issue 2

Full text available:  [pdf\(1.49 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We introduce the Iceberg-CUBE problem as a reformulation of the datacube (CUBE) problem. The Iceberg-CUBE problem is to compute only those group-by partitions with an aggregate value (e.g., count) above some minimum support threshold. The result of Iceberg-CUBE can be used (1) to answer group-by queries with a clause such as HAVING COUNT(\*) >= X, where X is greater than the threshold, (2) for mining multidimensional association rules, and (3) to complement existing strategies for identifying ...

#### 7 A powerful and SQL-compatible data model and query language for OLAP

Dennis Pedersen, Karsten Riis, Torben Bach Pedersen

January 2002 **Australian Computer Science Communications , Proceedings of the thirteenth Australasian conference on Database technologies - Volume 5**, Volume 24 Issue 2

Full text available:  [pdf\(1.12 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In this paper we present the SQLM OLAP data model, formal algebra, and query language that, unlike current OLAP data models and languages, are both *powerful*, meaning that they support irregular dimension hierarchies, automatic aggregation of data, and correct

aggregation of data, and *SQL-compatible*, allowing seamless integration with relational technology. We also consider the requirements to the data model posed by integration of OLAP data with external XML data. ...

**Keywords:** OLAP, data integration, data models, multidimensional databases, query languages

## 8 Visualization: Query, analysis, and visualization of hierarchically structured data using Polaris

Chris Stolte, Diane Tang, Pat Hanrahan

July 2002 **Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining**

Full text available:  [pdf\(10.02 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In the last several years, large OLAP databases have become common in a variety of applications such as corporate data warehouses and scientific computing. To support interactive analysis, many of these databases are augmented with hierarchical structures that provide meaningful levels of abstraction that can be leveraged by both the computer and analyst. This hierarchical structure generates many challenges and opportunities in the design of systems for the query, analysis, and visualization of ...

## 9 An overview of data warehousing and OLAP technology

Surajit Chaudhuri, Umeshwar Dayal

March 1997 **ACM SIGMOD Record**, Volume 26 Issue 1

Full text available:  [pdf\(101.60 KB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

Data warehousing and on-line analytical processing (OLAP) are essential elements of decision support, which has increasingly become a focus of the database industry. Many commercial products and services are now available, and all of the principal database management system vendors now have offerings in these areas. Decision support places some rather different requirements on database technology compared to traditional on-line transaction processing applications. This paper provides an overview ...

## 10 Tracking three dimensional moving light displays (abstract only)

Michael Jenkin

January 1984 **ACM SIGGRAPH Computer Graphics**, Volume 18 Issue 1

Full text available:  [pdf\(3.92 MB\)](#) Additional Information: [full citation](#), [abstract](#)

A method is presented for tracking the three-dimensional motion of points from their changing two-dimensional perspective images as viewed by a nonconvergent binocular vision system. The algorithm relies on a general smoothness assumption to guide the tracking process, and application of the tracking algorithm to a three-dimensional moving light display based on Cutting's Walker program as well as other domains are discussed. Evidence is presented relating the tracking algorithm to certain belief ...

## 11 The perception of coherent motion in two-dimensional patterns (abstract only)

Edward H. Adelson, J. Anthony Movshon

January 1984 **ACM SIGGRAPH Computer Graphics**, Volume 18 Issue 1

Full text available:  [pdf\(3.92 MB\)](#) Additional Information: [full citation](#), [abstract](#)

When one looks at a two-dimensional scene of moving objects, one can usually assign a velocity to each point in that scene with little effort. This suggests that some early visual processes are able to generate a two-dimensional velocity map using fast parallel computations. But it is not obvious how this should be done, and we are currently trying to

understand how the human visual system does it.

## 12 User-cognizant multidimensional analysis

Sunita Sarawagi

September 2001 **The VLDB Journal – The International Journal on Very Large Data Bases**, Volume 10 Issue 2-3

Full text available:  [pdf\(248.65 KB\)](#) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Our goal is to enhance multidimensional database systems with a suite of advanced operators to automate data analysis tasks that are currently handled through manual exploration. In this paper, we present a key component of our system that characterizes the information content of a cell based on a user's prior familiarity with the cube and provides a context-sensitive exploration of the cube. There are three main modules of this component. A Tracker, that continuously tracks the parts of the cub ...

**Keywords:** Maximum entropy, Multidimensional data exploration, OLAP, Personalized mining, User-sensitive interest measure

## 13 Query optimization by using derivability in a data warehouse environment

J. Albrecht, W. Hümmer, W. Lehner, L. Schlesinger


November 2000 **Proceedings of the 3rd ACM international workshop on Data warehousing and OLAP**

Full text available:  [pdf\(182.97 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

## 14 Research sessions: compression: Dwarf: shrinking the PetaCube

Yannis Sismanis, Antonios Deligiannakis, Nick Roussopoulos, Yannis Kotidis

June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data**

Full text available:  [pdf\(1.38 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Dwarf is a highly compressed structure for computing, storing, and querying data cubes. Dwarf identifies prefix and suffix structural redundancies and factors them out by coalescing their store. Prefix redundancy is high on dense areas of cubes but suffix redundancy is significantly higher for sparse areas. Putting the two together fuses the exponential sizes of high dimensional full cubes into a dramatically condensed data structure. The elimination of suffix redundancy has an equally dramatic ...

## 15 A case for dynamic view management

Yannis Kotidis, Nick Roussopoulos

December 2001 **ACM Transactions on Database Systems (TODS)**, Volume 26 Issue 4

Full text available:  [pdf\(892.57 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Materialized aggregate views represent a set of redundant entities in a data warehouse that are frequently used to accelerate On-Line Analytical Processing (OLAP). Due to the complex structure of the data warehouse and the different profiles of the users who submit queries, there is need for tools that will automate and ease the view selection and management processes. In this article we present DynaMat, a system that manages dynamic collections of materialized aggregate views in a data warehouse ...

**Keywords:** Data cube, OLAP, data warehousing, materialized views

**16 Understanding facts in a multidimensional object-oriented model**

Alberto Abelló, José Samos, Fèlix Saltor

November 2001 **Proceedings of the 4th ACM international workshop on Data warehousing and OLAP**Full text available:  [pdf\(1.37 MB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

"On-Line Analytical Processing" tools are used to extract information from the "Data Warehouse" in order to help in the decision making process. These tools are based on multidimensional concepts, i.e. facts and dimensions. In this paper we study the meaning of facts, and the dependencies in multidimensional data. This study is used to find relationships between cubes (in an Object-Oriented framework) and explain navigation operations.

**Keywords:** cube operations, functional dependencies, multidimensionality, object-oriented modeling

**17 XML and architecture: Achieving adaptivity for OLAP-XML federations**

Dennis Pedersen, Torben Bach Pedersen

November 2003 **Proceedings of the 6th ACM international workshop on Data warehousing and OLAP**Full text available:  [pdf\(119.23 KB\)](#)Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Motivated by the need for more flexible OLAP systems, this paper presents work on logical integration of external data in OLAP databases, carried out in cooperation between the Danish OLAP client vendor targit and Aalborg University. Flexibility is ensured by supporting XML as the external data format, since almost all data sources can be efficiently wrapped in XML. Earlier work has resulted in an extension of the targit system, allowing external XML data to be used as dimension ...

**Keywords:** OLAP, XML, adaptivity, federated databases

**18 Dissertation Abstracts in Computer Graphics**January 1992 **ACM SIGGRAPH Computer Graphics**, Volume 26 Issue 1Full text available:  [pdf\(2.53 MB\)](#)Additional Information: [full citation](#)**19 Perception of rotation in depth: the psychophysical evidence (abstract only)**

Myron L. Braunstein

January 1984 **ACM SIGGRAPH Computer Graphics**, Volume 18 Issue 1Full text available:  [pdf\(3.92 MB\)](#)Additional Information: [full citation](#), [abstract](#)

There are a variety of ways in which motion in the environment can provide information about three-dimensional relationships. One transformation that has received increasing attention in both the visual perception literature and in the machine vision literature is rotation in depth. This transformation, which includes any rigid rotation other than a rotation about the line of sight, can provide both a strong impression of depth and specific information about three-dimensional relationships in a ...

**20 Knowledge-based animation (abstract only)**

David Zeltzer

January 1984 **ACM SIGGRAPH Computer Graphics**, Volume 18 Issue 1Full text available:  [pdf\(3.92 MB\)](#)Additional Information: [full citation](#), [abstract](#)

In constructing a goal-directed system for automatic motion synthesis for computer animation, the essential problem is to account for the extraordinary flexibility and adaptability exhibited by moving creatures. The selective *potentiation* and *depotentiation* of elements of a hierarchy of motor control programs is a key to the generation of adaptive motor control. The constraints on motion sequences are analyzed, and mechanisms for achieving continuity of movements are discussed. The ...

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2004 ACM, Inc.  
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)